REMARKS

Claims 1-20 were pending in the patent application. By this Amendment, Applicants cancel Claims 2, 3, 14, 19 and 20 and add new Claims 21-26. The language of the newly-added claims is supported by the original Specification, for example in Fig. 4 and at page 16, line 4 through page 17, line 2. An additional filing fee of $1440.00 accompanies this submission, including the RCE filing fee of $810.00 plus $630.00 for the addition of 3 new independent claims, with the total number of claims remaining at 20 total. Should any additional fees be necessary, authorization is hereby given to charge Deposit Account 50-0510.

The Examiner had finally rejected Claims 1-20 under 35 USC 103 as unpatentable over Yang in view of Lim and further in view of Miller. Applicants respectfully contend that the invention is patentable over the cited art.

The present invention is directed to an apparatus, program storage device, and a method for evaluating workload across a processing environment having a plurality of computer systems each having a plurality of assigned workload units, wherein the method comprises the steps of assigning a plurality of impact values, one impact value for each workload unit assigned for each of the plurality of computing systems, wherein the assigning of each impact

YOR920000461**US1**   **-15-**

value comprises determining a change in system expiration date should the workload unit be removed from the system; and assessing the workload based on the impact values.

An impact value is assigned for each workload unit, wherein a workload unit is expressly defined for the application as "a subset of the workload", the workload being "the set of identifiable tasks that execute in the processing system" (see: page 8, line 19-page 9, line 3). For each subset of the workload, an impact value is assigned, representing the change in system expiration date that would occur if the workload unit was removed from the system. The term "expiration date" is the date when the server workload is expected to exceed its capacity because of growth in workload. The expiration date may be calculated using life expectancy, capacity space, or other method, as detailed in the Specification on page 11, lines 1-18.

Applicants expressly claim a method (independent Claims 1 and 21 and the claims which depend therefrom), apparatus (Claims 10, 11, 25 and 26) and a program storage device (Claims 12 and 23 and the claims which depend therefrom) for evaluating workload across a processing environment and for altering the workload based on impact value of workload units.

The Yang patent, is directed to a method and apparatus for modeling or profiling a system based on workload in

order design the system. A model is generated using "a set of generic system activities" and performance estimates, after which hardware parameters are determined to design a system to handle the activities. Yang uses computer activity elements ("CAE") "to capture the fundamental activities being performed by [a] respective software application" (see: Col. 4, line 1-8). A set of CAE is calculated for each transaction (see: Col. 26, lines 18-19). A "user focused workload" ("UFW") is collected "using forms or templates ...regarding the application that is available to the user" (Col. 5, lines 50-53). The UFW is translated into a computer activity workload ("CAW") representing the workload data structure that mathematically describes the workload in terms of CAE elements (Col. 5, lines 58-62). Yang models expected application workload for a user and then recommends system requirements for creating a system for that specific user for the user's expected usage of the system. As is clearly taught by Yang, the hardware requirements are determined based on modeled workload. Step 112 in Fig. 1 expressly shows that the final step of Yang's process is to determine hardware parameters. Similarly, step 410 of Fig. 4 shows a final step to "profile hardware system" and step 510 in Fig. 5 shows the final step to "determine possible hardware configurations". Yang in not analyzing an existing processing environment with existing

computer systems having assigned workloads and trying to determine how to re-distribute workload among them. Rather, Yang is modeling anticipated workload and designing a system to accommodate that workload. Yang does not determine the impact value of existing workload units in order to juggle those workload units among existing computer systems. Yang determines anticipated CAEs and then recommends an appropriate amount of hardware resources to accommodate the number of CAEs workload.

Applicants respectfully assert that the Yang patent neither teaches nor suggests the invention as claimed. With specific reference to the language of the independent claims 1, 10, 12, 21, 23, and 25, Yang does not teach or suggest evaluating workload across an actual processing environment having a plurality of computer systems each having a plurality of assigned workload units. Rather, Yang models anticipated requirements for projected applications under user-estimated usage conditions. The Examiner has, in the **Response to Arguments** section, stated that the recitation of "a processing environment having a plurality of computer systems each having a plurality of assigned workload units" has not been given patentable weight because the language appears in the preamble. Applicants respectfully point out that the same language appears in the body of the claims and that the claimed means and process steps cannot "stand

YOR920000461**US1**   **-18-**

alone" without that language. The claims clearly recite assigning impact values to workload units on each of the plurality of computer systems in a processing environment, assessing the impact of moving workload units from one computer system to another computer system and altering the workload among at least two of the computer systems. Clearly, the body of the claim depends on the existence of a computing environment comprising a plurality of computer systems.

In rejecting the steps and means for assigning a plurality of impact values to assigned workload units (Claims 1, 10, 12 and 22, 24 and 26), the Examiner generally cites the teachings from Col. 5, line 1 to Col. 6, line 19. Applicants respectfully assert that the cited passage does not teach or suggest the assigning of impact values. Yang generates a workload data structure to mathematically describe workload estimates by assigning CAEs in a processing "vacuum", such that an existing processing environment and existing computer system's workload is not taken into account. In contrast, and as specifically taught by the present Specification at page 5, line 7-9, the measure/impact value of a workload unit "is calculated in relation to a specific processing system and the other workload assigned to the system". Where the current claims expressly recite the assignment of impact values relative to

the existing workload and system expiration date, Yang uses theoretical values for defining CAEs and the amount of resources required for CAEs (see: the cited passage from Col. 5, line 1 to Col. 6, line 19). While Yang may use prior/historical averages of resource requirements based on data from other systems (Col. 6, lines 9-12), Yang is neither teaching nor suggesting determining impact values for an existing system having an assigned workload.

Further, Yang does not teach or suggest that a system has an expiration date, since Yang does not have a predefined, existing processing environment having a plurality of computer systems each having a plurality of assigned workload units. Rather, Yang models a larger system to accommodate larger workloads, so the concept of system expiration date is simply not applicable. Applicants note that the Examiner has not cited any teachings in Yang against the "system expiration" language.

Applicants reiterate that the terms used in the claims, including "impact value" and "system expiration date" are terms that have definite meanings for the present invention. The Yang patent does not teach or suggest those terms for its theoretical, simulated environment, let alone in the context of workload evaluation for an existing processing environment having a plurality of computer systems each having assigned workload units.

YOR920000461US1   -20-

In rejecting the independent claim features of the steps and means for assessing the impact of moving a workload unit from a donor computer system to a recipient computer system based on impact values (Claims 1, 10 and 12 and Claims 21-26), the Examiner has cited the passage from Col. 33, lines 30-62. What Yang teaches in the cited passage is that different possible hardware configurations are considered for a modeled workload. Yang looks at the resources of a hardware configuration and determines whether the resources are sufficient to handle the modeled CAW. If not, Yang looks at a different projected hardware configuration to handle the CAW. If more than one hardware configuration is suitable, Yang may further look at other "desirability" factors (e.g., cost) in selecting a hardware configuration. Yang neither teaches nor suggests assessing the impact of moving some of the CAEs of the CAW to a different hardware configuration.

The Examiner has acknowledged that Yang does not teach a processing environment having a plurality of computer systems each having a plurality of assigned workload units. The Examiner has cited the Lim patent for teaching distributing workload from a system (server) among other servers. Applicants respectfully assert that one having skill in the art would not be motivated to modify Yang with Lim, since Yang is modeling systems for system design while

Lim is directed to workload balancing among servers. Furthermore, even if one skilled in the art was motivated to modify Yang with Lim, one would not arrive at the present invention since neither Yang nor Lim teaches steps, program storage device, or means for assigning a plurality of impact values, one impact value for each actual workload unit assigned for each of the plurality of computer systems in a processing environment having a plurality of computing systems each having a plurality of assigned workload units, does not teach assigning impact values by determining a change by assigning an impact number representing the number of days that the expiration date of the computer system would be changed if a workload unit is removed from the system with all other workload units remaining the same; and does not teach assessing the impact of moving the workload from a donor computer system to a recipient computer system based on said impact values. As discussed above, Yang models generic system activities that are "independent of any particular hardware system" (Col. 4, lines 1-3).

Lim teaches a name service system that distributes its work among distributed servers. Lim exchanges messages with servers to determine server availability. When a server does not send a reply message or if a "connect" fails (Col. 2, lines 49-51 and Col. 9, lines 36-52), it is assumed that the server is not available. In addition, Lim expressly

teaches that servers whose workload exceeds a predetermined workload value are deleted from the list of available servers (Col. 2, lines 41-43). Applicants respectfully contend that Lim does not teach or suggest those claim features which are missing from the Yang patent. Lim does not teach or suggest balancing workload among servers, Lim just sends work to any available servers to which connections can be made. Further, Lim does not teach or suggest assigning impact values or assessing the impact of moving workload from one server to another. Lim simply removes a server from the list of available servers if its workload exceeds a predetermined workload value. Applicants respectfully assert that the Examiner erred in concluding that "it would be obvious the capacity planning of Yang could be implemented with the workload balancing among a plurality of servers as in Lim". The capacity planning of Yang simply says "add more hardware" or "add more servers" in a Lim system. Since neither Yang nor Lim teaches assessing the impact of moving workload from one entity to another, the combination of Yang and Lim does not obviate the invention as claimed.

In acknowledging that Yang and Lim fail to teach assigning an impact number representing the number of days that the expiration data of the computer system would be changed with all other workload units remaining the same,

the Examiner has newly cited the Miller patent. Miller teaches modifying tasks, themselves (Col. 3, lines 29-37), to require less computing capability. Miller does not teach or suggest assigning impact values to workload units as they relate to system expiration date and then assessing the effect of moving workload units from one computer system to another. Rather, Miller expressly teaches modifying the workload units, themselves, (see: Fig. 2, step 52) so that the workload units require less processing time or fewer resources. The cited passage from Col. 9, lines 3-68 teach the initial Miller assessment of task requirements and resource availability, but do not teach or suggest assigning impact values related to system expiration date. Further, since Miller teaches away from moving workload among resources and instead teaches modifying tasks at assigned resources, it cannot be concluded that Miller provides the teachings which are missing from the combination of Yang and Lim.

For a determination of obviousness, the prior art must teach or suggest all of the claim limitations. "All words in a claim must be considered in judging the patentability of that claim against the prior art" (In re Wilson, 424 F. 2d 1382, 1385, 165 U.S.P.Q. 494, 496 (C.C.P.A. 1970). Since the cited Yang and Lim references fail to teach each and

every one of the claim limitations, a *prima facie* case of obviousness has not been established by the Examiner.

With specific reference to the language of the dependent claims, Applicants note that the Examiner has repeated the previous rejections and citations. Applicants, therefore, maintain the contention that the combination of references does not obviate the claims.

With regard to Claims 4 and 13, the Examiner has again cited the passage, and code detailed therein, from Col. 26, line 45-Col. 27, line 5. Yang detailed calculating anticipated reads and writes for one transaction to model the workload for the transaction. Such is not the same as or suggestive of sorting actual workload units based on assigned impact values.

With regard to the language of Claims 5 and 14, which has been added to the independent claims by this amendment, the Examiner again cited the passage from Col. 25, lines 13-20 which teaches characterizing one transaction in terms of CAE. There is not mention in the passage of altering workload of one or more system. Applicants reiterate that since Yang is "operating" in a theoretical environment, Yang neither teaches nor suggests actual expiration dates of actual systems or altering actual workloads for one or more systems.

With respect to Claims 6 and 15, and Claims 7-9 and 15-18 which depend respectively therefrom, Applicants argue that the cited passage from Col. 33, lines 30-62 does not anticipate or obviate the claimed step and means for comparing the expiration data of each of a plurality of actual computing systems to at least one target planning date for servicing the actual computing systems. What Yang mentions in the cited passage is response times for modeled systems. Again, since Yang is operating in the theoretical environment, real factors such as servicing dates are simply not relevant. Yang does not teach or suggest any consideration of target planning dates for servicing systems. Clearly, therefore, Yang does not teach or suggest the language of Claims 6-9 and 15-18.

In rejecting the language of Claims 7 and 16, the Examiner cites Col. 25, lines 13-20 of Yang. The cited passage states that "the transaction mix can be altered by changing the setting for number of transactions, the frequency, or the percentage mix for each transaction type." Those teachings relate to Yang generating a set of CAE for (i.e., modeling) a transaction. There is nothing in the cited passage which teaches or suggests expiration dates, target planning dates for system servicing, or altering workloads for at least two actual computing systems.

With regard to the language of Claims 8 and 17, the Examiner again cites the Yang teachings found in Col. 6, lines 9-36. The cited passage states that workload definition information for use in modeling can be collected from actual workload information, that the workload definition is "transferable to other applications and hardware", that "the same workload definition can be used for analysis performed with respect to a plurality of different hardware platforms". Creating a generic workload definition for system modeling does not anticipate the claims which recite creating "From" and "To" lists relative to at least one target planning date for servicing of actual computing systems, and reassigning actual workload units based on assigned impact values reflecting a change in system expiration date if a workload unit is moved from one actual computing system to another. Yang is providing a generic workload definition that can be used to evaluate different candidate system. Such is not the same as or suggestive of evaluating and moving actual workload units in a runtime environment.

The language of Claims 9 and 18, which recites calculating new expiration dates for computing systems on the "From" and "To" lists after reassignment of workload units, has additionally been rejected. The Examiner cites Col. 5, line 1 through Col. 6, line 36 of Yang. The cited
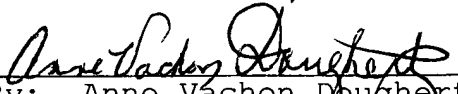
teachings all relate to how Yang defines its theoretical workload and CAE. Applicants again argue that the Yang teachings do not anticipate system expiration dates at all, since it is modeling candidate systems, and do not anticipate calculating new expiration dates of actual systems as claimed. Applicants reiterate that Yang is essentially designing a system and opts to change the system design, but does not teach of suggest altering the workload of an existing system.

With regard to Claim 11, Applicants acknowledge that Fig. 7 of Yang illustrates a storage location. However, Yang's provision of a storage location in a "mechanism for profiling a system" is not the same as or suggestive of providing a storage location in the apparatus of Claim 10, including an administrative processor comprisingan impact value component for assigning a plurality of impact values, one impact value for each workload unit assigned to each of the plurality of computing systems in a processing environment having a plurality of computer systems each having a plurality of assigned workload units, wherein said assigning of each impact value comprises determining a change in system expiration date should a workload unit be removed from the system; and a processing component for assessing the workload based on said impact values.

Appellants reiterate that the prior art must teach or suggest all of the claim limitations. "All words in a claim must be considered in judging the patentability of that claim against the prior art" (In re Wilson, 424 F. 2d 1382, 1385, 165 U.S.P.Q. 494, 496 (C.C.P.A. 1970). Since the cited Yang and Lim references fail to teach each and every one of the claim limitations, a *prima facie* case of obviousness has not been established by the Examiner. Further, obviousness cannot be maintained without some teaching or suggestion of the claim features. The Federal Circuit has stated that when patentability turns on the question of obviousness, the obviousness determination "must be based on objective evidence of record" and that "this precedent has been reinforced in myriad decisions, and cannot be dispensed with." (In re Lee, 277 F. 3d 1338, 1343 (Fed. Cir. 2002)). Moreover, the Federal Circuit has stated that "conclusory statements" by an examiner fail to adequately address the factual question of motivation, which is material to patentability and cannot be resolved "on subjective belief and unknown authority" (Id. at 1343-1344). Applicants respectfully assert that the Examiner has not cited actual teachings but has provided conclusory statements about the teachings of Yang, Lim and Miller without basis in actual teachings from those patents.

Based on the foregoing amendments and remarks, Applicants respectfully request entry of the amendments, reconsideration of the amended claim language in light of the remarks, withdrawal of the rejections, and allowance of the claims.

Respectfully submitted,
L. B. Mummert, et al

By: Anne Vachon Dougherty
Registration No. 30,374
Tel. (914) 962-5910